



# Cloud Computing

**Bill Vass**  
**President/COO**  
**Sun Microsystems Federal, Inc.**

**“Recovering CIO”**  
[bill.vass@sun.com](mailto:bill.vass@sun.com)  
[blogs.sun.com/BVass](http://blogs.sun.com/BVass)

# Everyone is Talking About Cloud Computing

Software as a Service

Platform as a Service

Storage as a Service

Grid Computing

Database as a Service

Virtualization

Utility Computing

Application Hosting

Infrastructure as a Service

# IDC on Cloud Computing:

***“This is about the IT industry’s new model for the next 20 years.”***

***Vernon Turner, head of enterprise infrastructure, consumer and telecoms research.***

**“...reliable services delivered through data centers ... accessible anywhere that has access to networking infrastructure.**

**The Cloud appears as a single point of access for all the computing needs...”**

# It's Not Just About Cheap Computing

Efficiency

Agility

# Efficiency

## Economics



- Pay as-you-go
- Op-ex vs. Cap-ex
  - SLA
- Virtualization

## Developer Centric



- Rapid,  
self provisioning
- Faster deployment
  - Self service
  - API-driven

## Flexibility



- Standard services
  - Elastic
  - On demand
  - Multi-tenant

# Changing IT Relationships

 <h2>Developers</h2>	 <h2>Deployers</h2>
<ul style="list-style-type: none"> <li>• Why won't IT support this?</li> <li>• Why can't I use the versions I want?</li> <li>• Why can't I get better availability?</li> <li>• How can I pay for what I need?</li> <li>• How quickly can I get more servers?</li> </ul>	<ul style="list-style-type: none"> <li>• Why do we have so many versions of everything?</li> <li>• Where can I cut costs?</li> <li>• How can I do finer grain provisioning?</li> <li>• Where do we enforce security, regulation and audit?</li> </ul>

# Driving IT Agility

## Current State

- Commercial software packages
- Relational data
- Management
- Server-scale

## Trend

- Open source communities
- Unstructured, rich data
- Analytics
- Network data center-scale

# All Clouds Share Key Traits

One Service Fits All

Virtualized Physical Resources

Self Provisioning

Elasticity

Pay Per Use

Programmatic Control

A photograph of a goat standing on a rocky cliff edge, looking out over a mountain range. The goat is brown and has small horns. The background shows a vast mountain landscape under a cloudy sky.

# Layers

---

## Business Models

---

## Application Domains

**BUT**

clouds  
can also  
be quite  
different



# Cloud Dynamics – Evolves like the Network & Internet

- **Just like the network,** will get availability through redundancy and multiple providers
- **Just like the network,** will need common open standards and open source to provide redundancy to meet SLAs
- **Just like the network,** will have:
  - Private (LAN)
  - Mixed (WAN)
  - Public (Internet / DMZ)
- **Just like the network,** will need encryption and trust to enable it
- **Just like the network,** will need monitoring and security management

# Types of Cloud Computing

## Cloud Storage

- Encrypted on the way out, Decrypted on the way in.
- Many different SLAs and types

## Application Cloud

- Encrypted communication
- Application server threads
- Stateless (most common) and Stateful

## Infrastructure VM Cloud

- Encrypted Communication
- Windows, Linux, Solaris OS Images
- Full Applications, Virtual Servers, Virtual Desktops (Stateless clients)

# Cloud Computing Layers

## Software as a Service

Applications offered on-demand over the network (salesforce.com)

## Platform as a Service

Developer platform with built-in services (Google App Engine, Microsoft Azure Platform)

## Infrastructure as a Service

Basic storage and compute capabilities offered as a service (Amazon web services, Microsoft's Cloud Infrastructure Services, Mosso)

# Business Models

## Public



You don't know who else is on the same server, network or disk that you are

## Private



You own the server, network and disk, and decide who gets to run on it with you

## Hybrid



You own some parts and are sharing some parts, though in a controlled way

# Application Domains

HPC

Medical

Intelligence

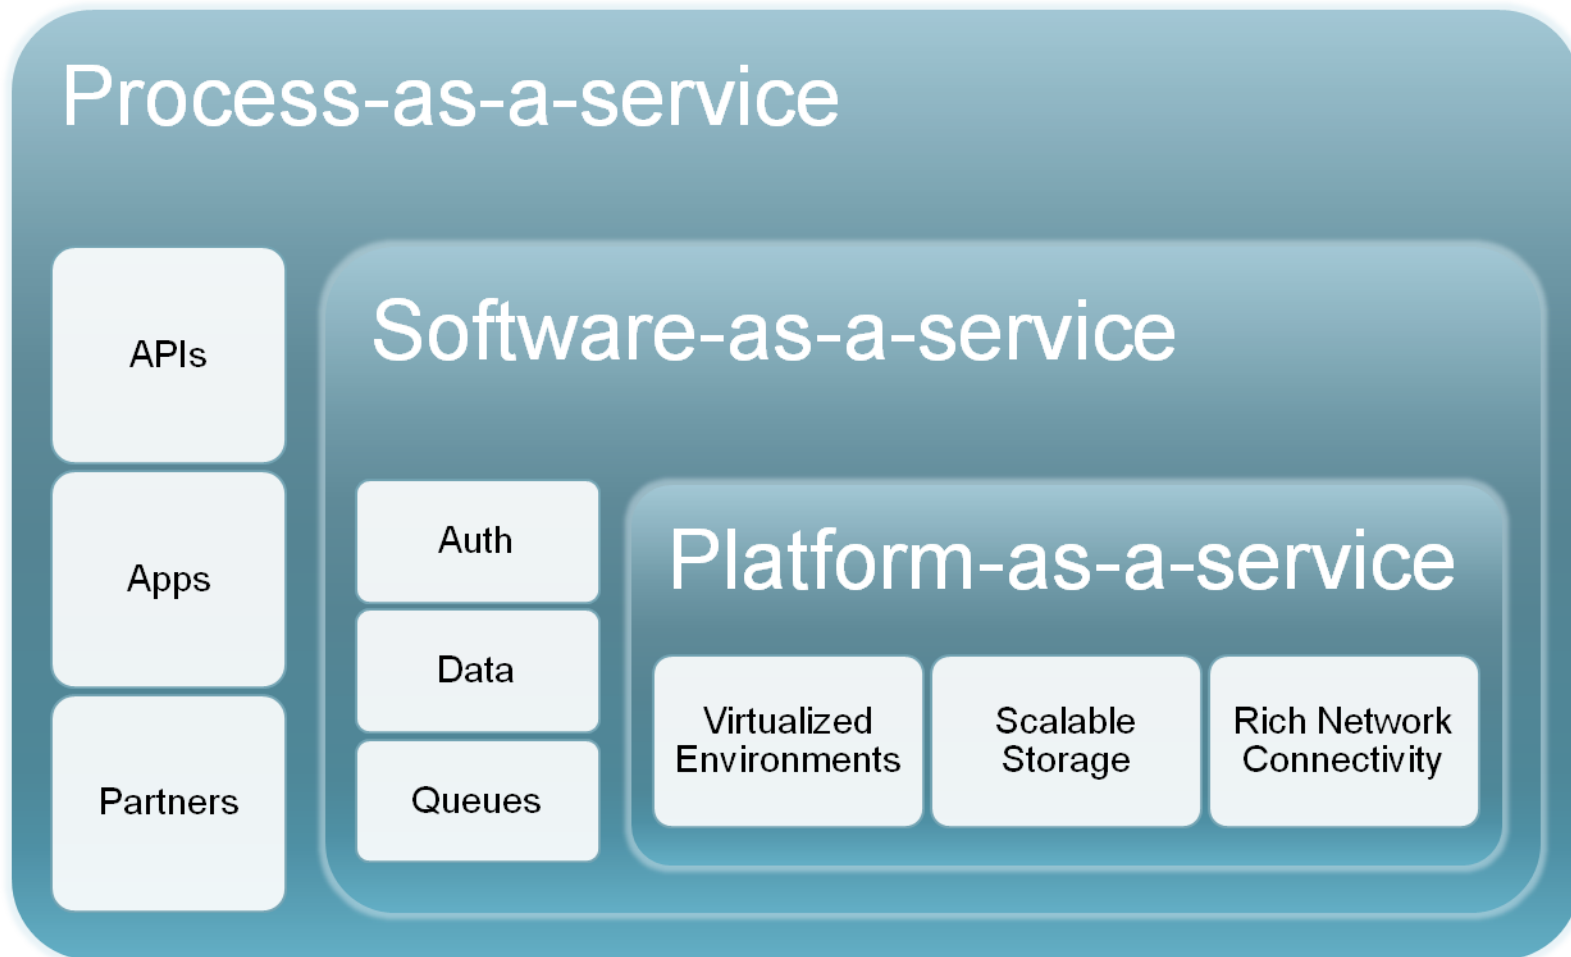
Finance

Analytics

Web

Domains Drive Differences  
in Hardware and Software Architecture

# The Cloud - Under the Hood



# Cloud Platforms, Compared

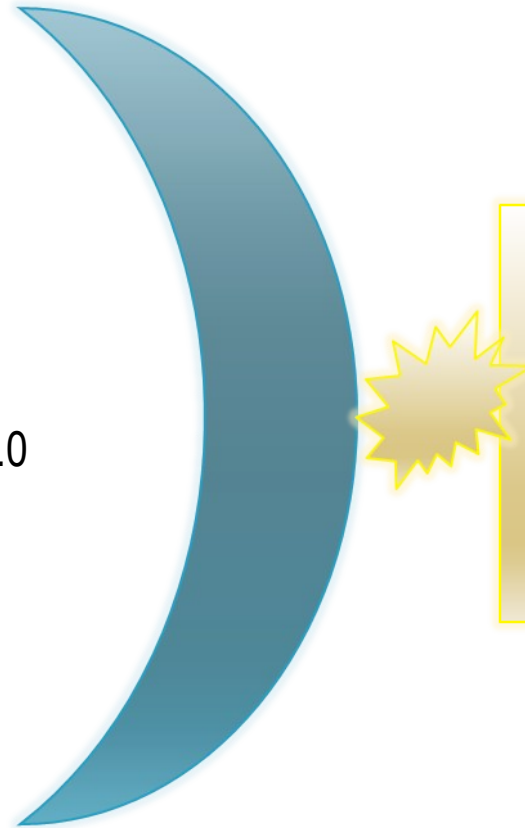
	AWS	AppEngine	Sun Cloud
Storage	S3, EBS	SimpleStorageAPI	Luster
Database	SimpleDB	CouchDB	MySQL
Queue	SQS	--	JavaSMQ
Virtualization	EC2 (Xen)	--	xVM OpsCenter
Framework	--	Django	JCAPS
User Accounts	--	Google Accounts	IDM Open SSO
Search	--	Google Search	SOLR+Lucene
Desktop	--	GoogleDocs?	SSR, VB, ODF
CDN/Cache	CloudFront		Varnish



# Cloud Dynamics

# Cloud as “Bow-Wave of Change”

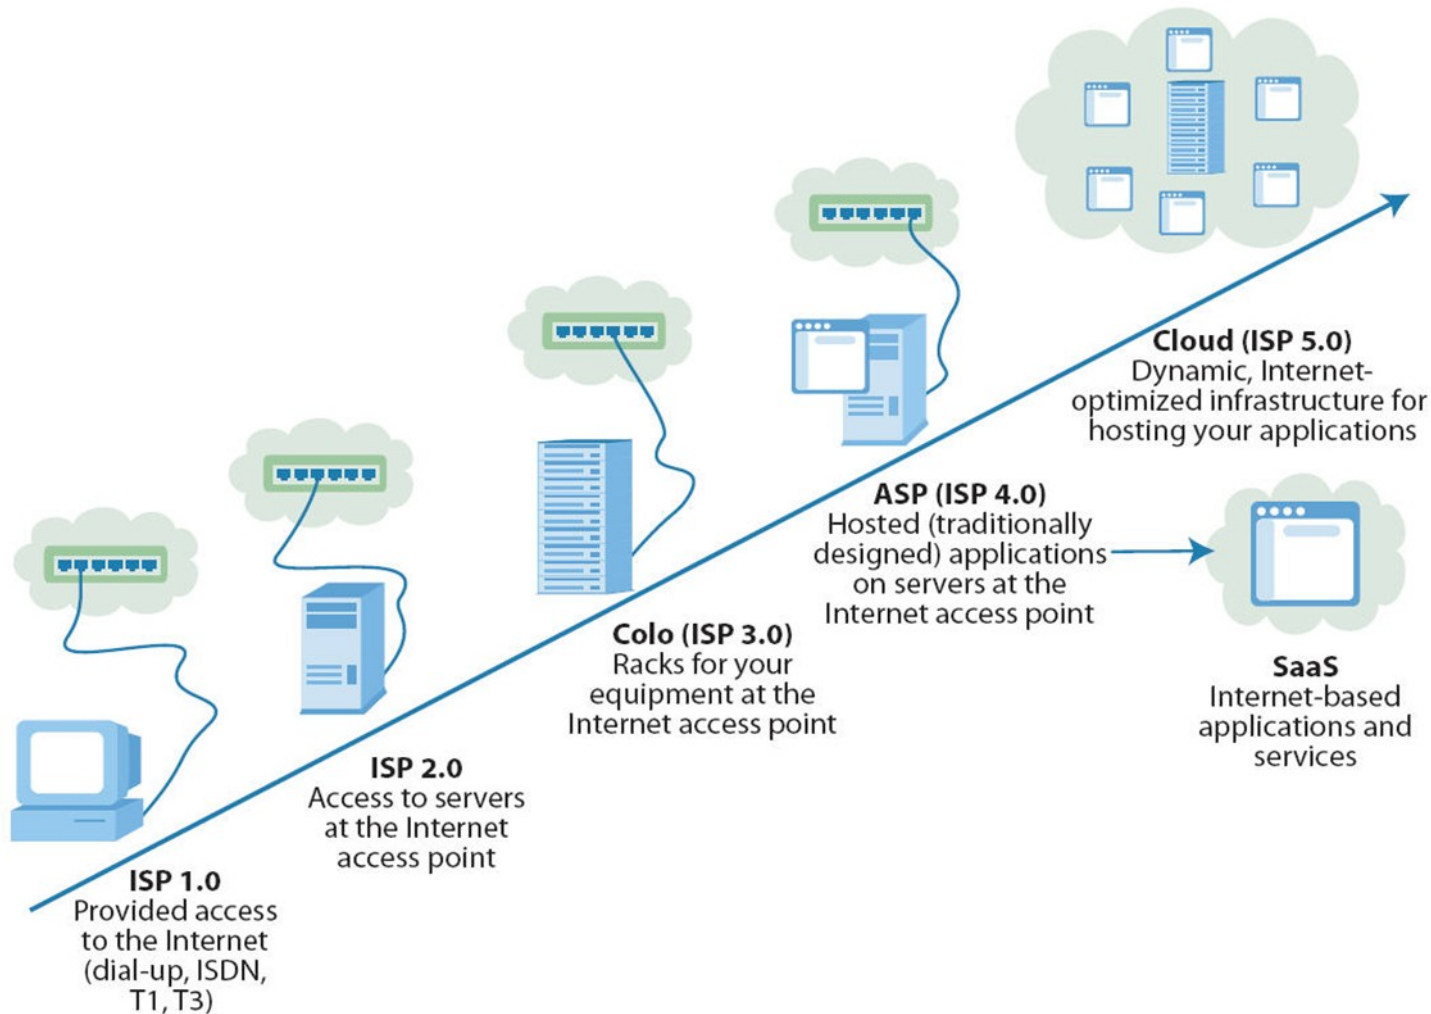
Virtualization  
 Social Networks  
 Mashups, DIY Tools  
 Exponential Change  
 Tech-savvy customers  
 Tech-savvy employees  
 Consumerization / Web 2.0  
 Cloud Computing  
 Public Infrastructure  
 Global Collaboration  
 Cultural Change  
 Security  
 Mobility  
 Green



- Employees
- Missions / Data Products
- Business Operations
- Partners / Collaborators
- The Public

Source: Doug Neal, CSC

# Cloud as “Evolution of Hosting”



Source: Forrester Research

# Changing Software Economics



# New Data and Management Economics

## Compute Trend

**New Analytics Emerge**  
(MapReduce, Hadoop...)



Architectural shift to the cloud  
and HPC-style workloads



**Greenplum**  
Open source, general  
purpose datawarehouse

**TERADATA.**

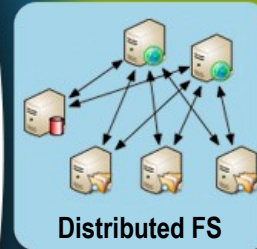
Proprietary, dedicated  
datawarehouse

**ORACLE**

OLTP is the  
datawarehouse

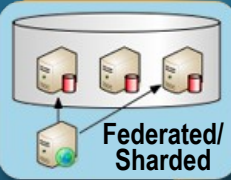
## Data (Storage) Trend

**Semi-structured Data**  
(Mogile, Bigtable, HDFS...)



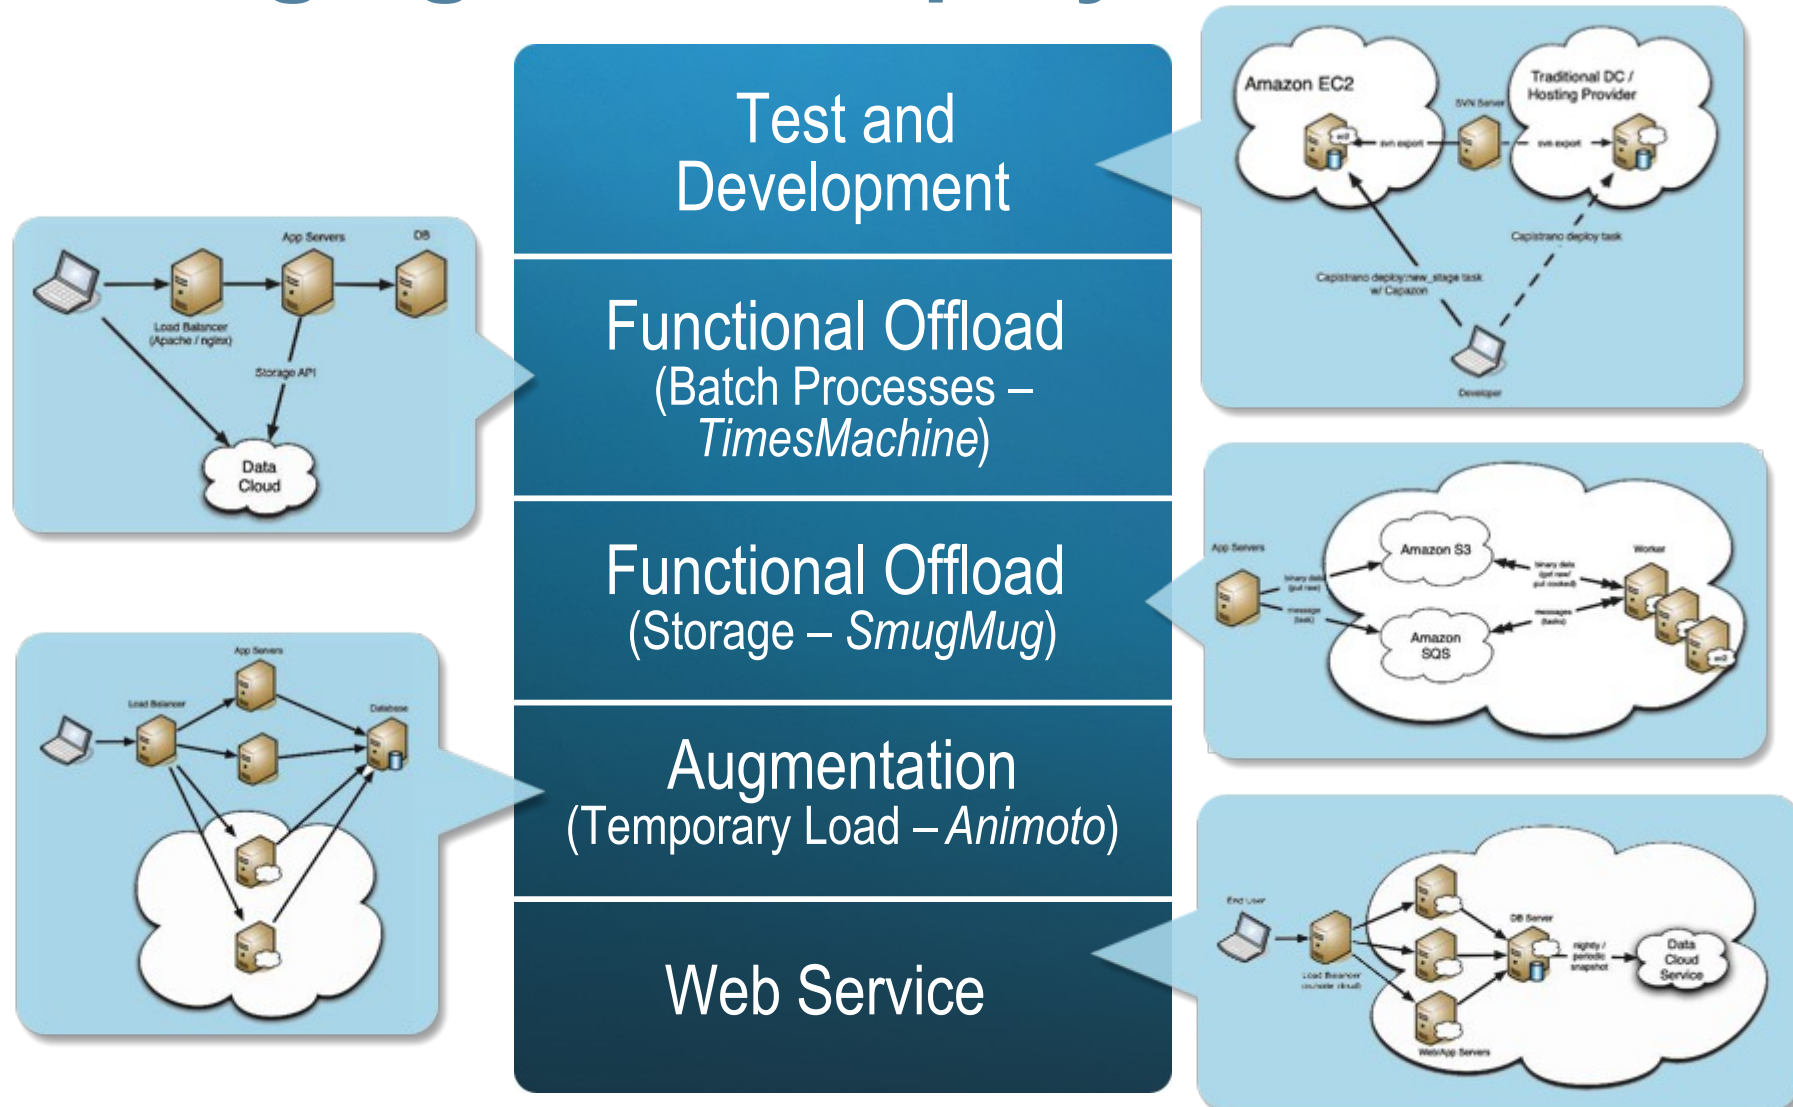
**Unstructured  
Data**

**Semi-structured  
Database**  
*ScaleDB, Big Table,  
SimpleDB hBase*







**Structured  
Data**



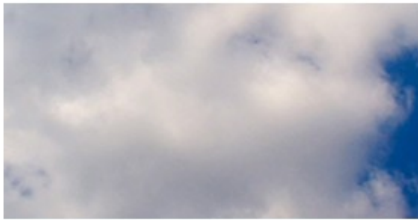

# Emerging Cloud Deployment Patterns



# Cloud Ownership Models

<b>USE</b> the Cloud <i>(no datacenter  ownerships)</i>	<b>LEVERAGE</b> the Cloud	<b>BUILD</b> My Own Internal Cloud	<b>BE</b> the Cloud
			
<ul style="list-style-type: none"> <li>• Startup</li> <li>• SMB</li> <li>• Research projects</li> </ul>	<ul style="list-style-type: none"> <li>• Temporary on-demand load</li> <li>• Functional off- load</li> </ul>	<ul style="list-style-type: none"> <li>• Enterprise infrastructure grid</li> <li>• Drive internal IT economics</li> <li>• Standardized development environment/ services</li> </ul>	<ul style="list-style-type: none"> <li>• Redefine services</li> <li>• New business offerings</li> <li>• Hosting and operations partners</li> <li>• Software vendors</li> </ul>

# Concerns of Cloud Layers

IaaS	PaaS	SaaS	Developers/ Deployer Tradeoffs
			
<ul style="list-style-type: none"> <li>• API-driven provisioning</li> </ul>	<ul style="list-style-type: none"> <li>• Reliability, load balancing, security subsumed</li> <li>• Data privacy, retention, recovery aspects</li> </ul>	<ul style="list-style-type: none"> <li>• User experience dominates</li> <li>• Loose coupling of other services</li> </ul>	<ul style="list-style-type: none"> <li>• Tradeoffs at each layer</li> </ul>

# Evolving Service Levels

## Reliability, Availability, Serviceability (RAS)

- Focused on state recovery
- Low slew rate of demand or change

## Network Scale Service Level Demands (PIPE)

- Predictability: SLAs under high rate of change
- Integrity: Security, compliance, correctness
- Productivity: Agility and time to service
- Efficiency: Throughput/\$, work/watt



# Examples

Results 1 - 10 of about 19,500,000 for [cloud computing](#).

## [Cloud computing security: Who knew?](#)

Computerworld - 11 hours ago

By Scott Bradner April 27, 2009 (Network World)  
less than perfect agreement on just what



## [DMTF Group to Create Cloud Computing Spec](#)

eWeek - 13 hours ago

The Distributed Management Task Force is creating a group to develop management standards for **cloud computing**. ...

[Open Standards For The Cloud](#) SYS-CON Media (press release)  
DMTF creates open cloud incubator. ITworld.com

## [SAS to build USD 70 million cloud computing facility](#)

Al-Bawaba - Apr 26, 2009

SAS, the leading provider of business intelligence software, will build a USD 70 million cloud computing facility in North Carolina. Demand source: \$70 million



## [Hosting.com Partners with VMware to Deliver Cloud Trial C](#)

HostSearch.com - Apr 26, 2009

As part of the VMware vCloud initiative, developers, IT professionals, SI are invited to experience **Cloud Computing** firsthand with 30 days ...

[Operating system for building the internal cloud](#) Help Net Security  
EMC integrates with VMware vSphere 4 for next generation. AME 1-1

## [Fourteen universities to study cloud computing with \\$5M NSF grant](#)

Geek.com - 10 hours ago

The program works with IBM and Google to look at the infrastructure requirements to

## [Fujitsu Launches Cloud Services](#)

Japan Corporate News (press release) - Apr 26, 2009

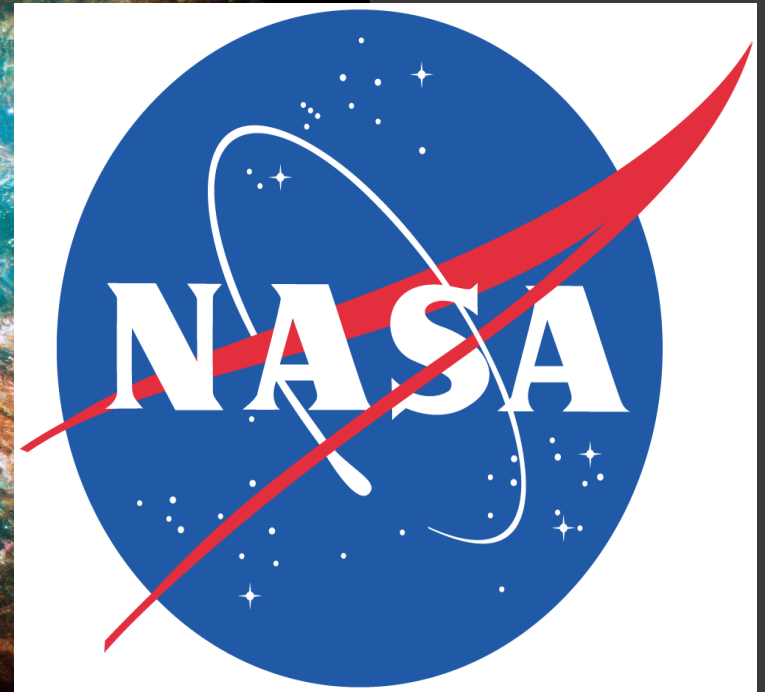
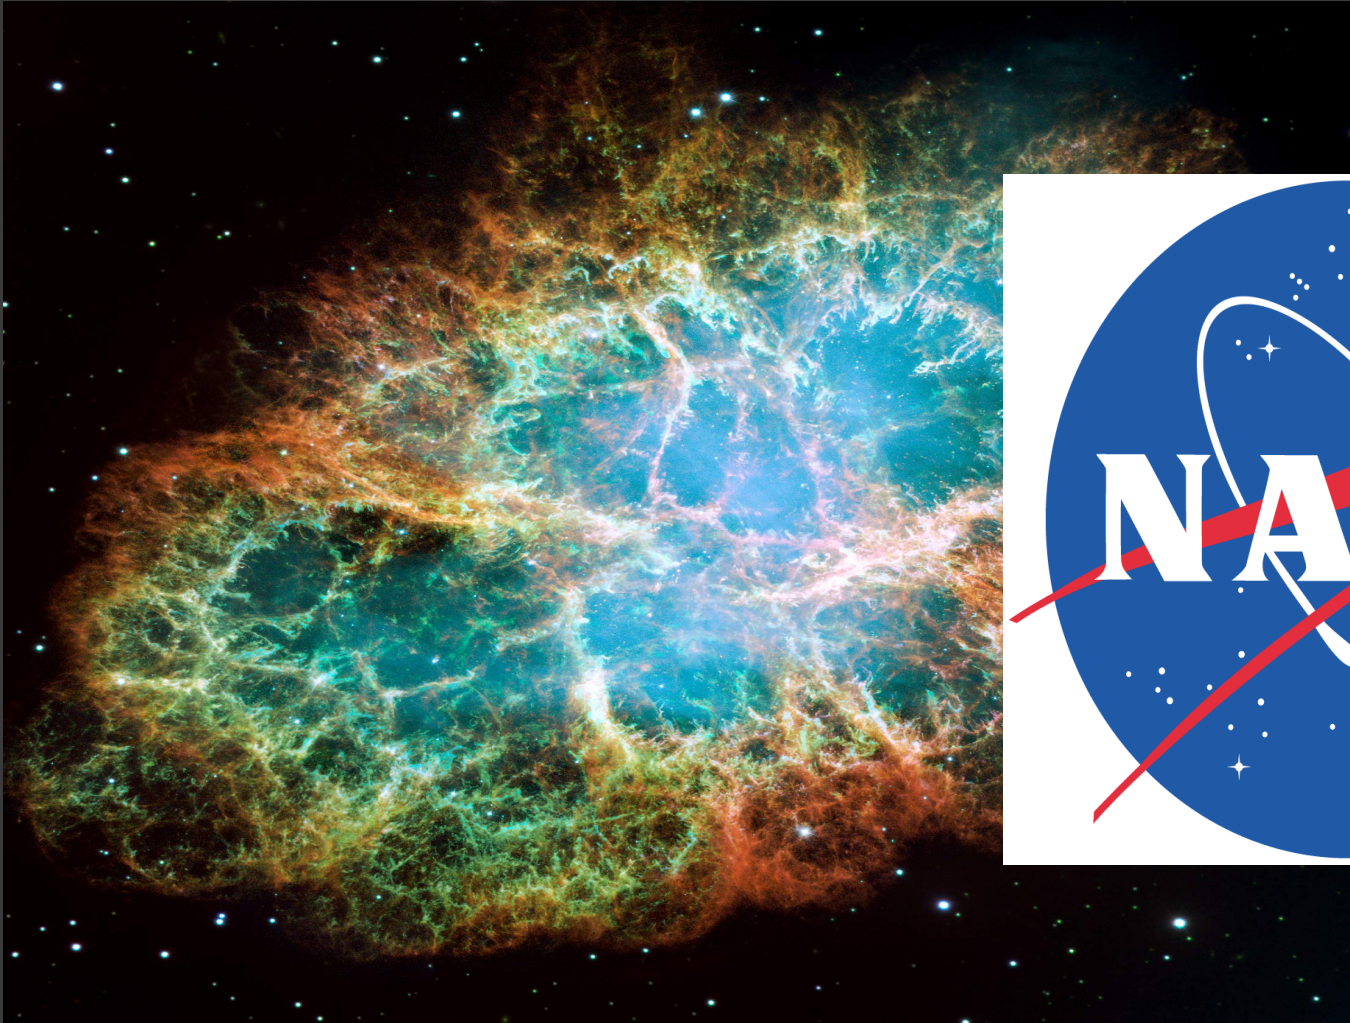
Fujitsu is seeking to address this issue by providing its Trustee has the following features. ...

## [Heading Back to Cloud Co](#)

SYS-CON Media (press release) - 2 hours ago

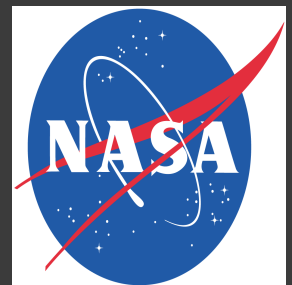
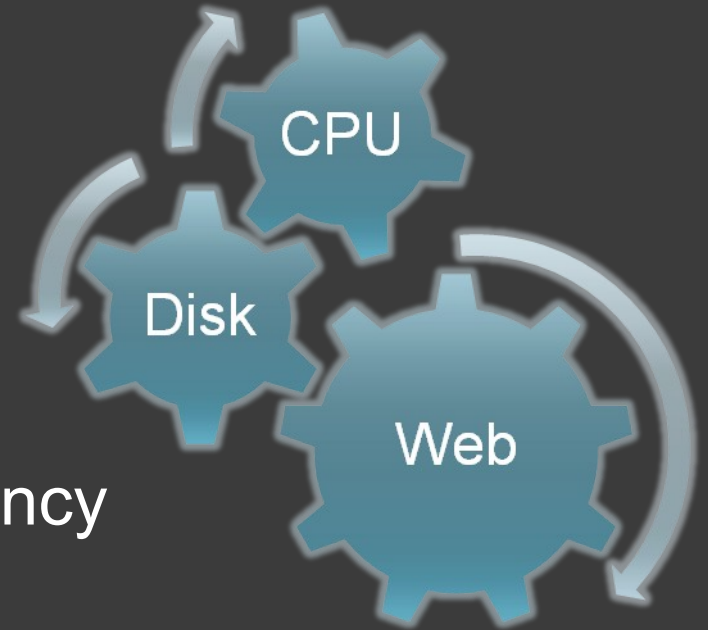
The merger does make Sun very relevant again, as its hardware is in the coming **Cloud Computing** deployment battle. ...

# Introducing NASA NEBULA



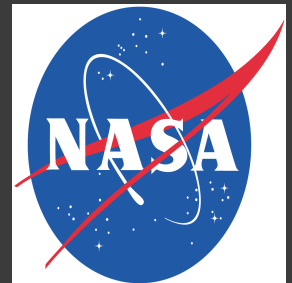
# NEBULA – Built for Science

- Science-Class Cloud Computing
  - High CPU-to-Disk Ratio
- Built for Research
  - Fastest networks in the Agency
  - Massively Parallel, Loosely Coupled



# NEBULA – Built for Collaboration

- ⦿ True Single-Sign-On, for the Public
- ⦿ Enterprise Search, across the Cloud
- ⦿ All Cloud apps live at [apps.nasa.gov](https://apps.nasa.gov)



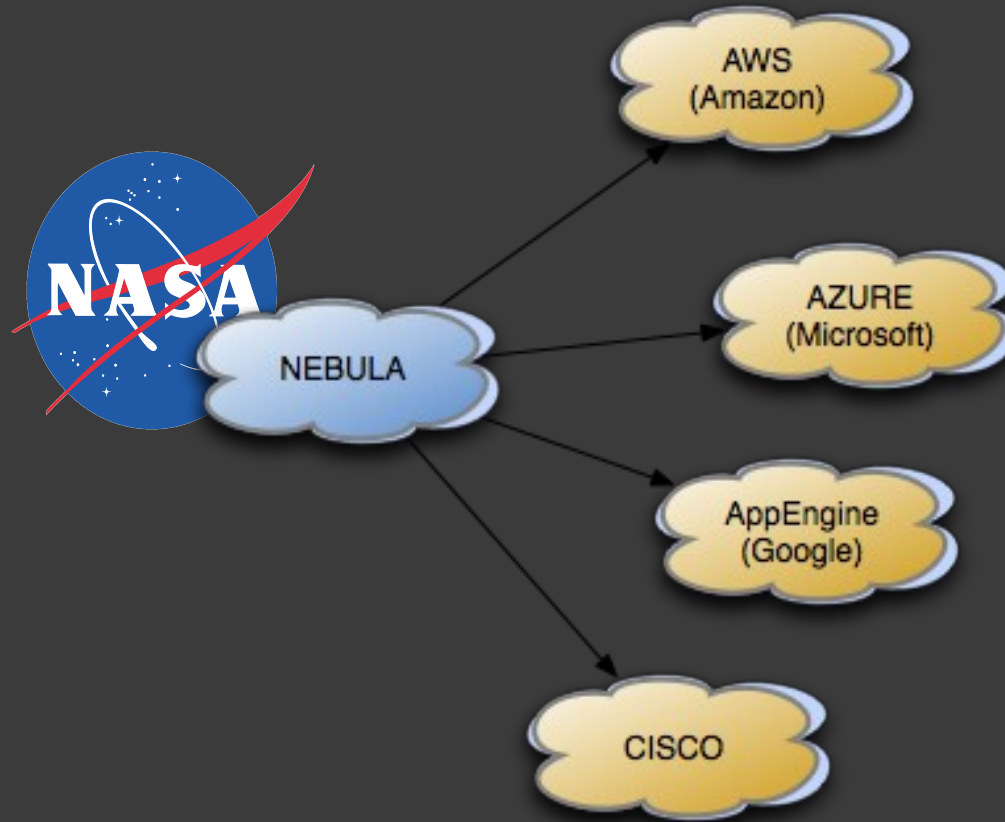
# NEBULA – Built for the Web

- ⦿ Friendly URLs
- ⦿ Designed for Search Engines, RSS, and aggregation
- ⦿ It's RSSable, Tweetable
- ⦿ What if NASA was on the first page of Google results for the term 'Space'?

# NEBULA – Built for Partners

- ◎ Your science partners can instantly connect from your NEBULA app, to their own research tasks within public Cloud Services (EC2, Azure, AppEngine)
- ◎ Your private fleet of Post-Docs can work on your data – at 10 cents an hour

# NEBULA – Built for Partners



# NEBULA – Built for Government

- ⦿ Policy compliant for contributions
- ⦿ Consolidated moderation interface
- ⦿ Everything-compliant (PII, First Amendment, COPPA, Section 508, etc)



# What *is* NEBULA?



The diagram illustrates the NEBULA services architecture. It features a large teal rectangle divided into five sections. The top section is a single wide bar labeled 'Services'. Below this, there are four vertical columns, each labeled with a service type: 'Software', 'Hardware', 'Networks', and 'Workflow'. These columns are separated by thin white lines. At the bottom of the diagram is a solid teal horizontal bar that spans the entire width of the four columns, representing a common base or platform.

## Services

Software

Hardware

Networks

Workflow

# Cloud Platforms, Compared

	AWS	AppEngine	NASA Nebula
Storage	S3, EBS	SimpleStorage API	Luster
Database	SimpleDB	CouchDB	MySQL
Queue	SQS	--	RabbitMQ
Virtualization	EC2 (Xen)	--	Eucalyptus
Framework	--	Django	Django
User Accounts	--	Google Accounts	eAuthentication
Search	--	Google Search	SOLR+Lucene
Desktop	--	GoogleDocs?	--
CDN/Cache	CloudFront		Varnish

# Application Time to Market

Current Web App Process	NEBULA Cloud Platform
Procure Server – 6-12 weeks	Procure VM – 60-120 seconds
Configure Server – 2-3 days	Included.
AWRS Filing – 2-3 days	Already done.
Set up Source Control – 2-3	Included.
Security Plan – 3 weeks, min.	Included.
SSL Certificates – 2-4 weeks.	Included.
Develop Terms of Use – 6	Included, for most collaboration.
eAuth Integration – 40 hours.	Included.
Develop Processes – 3-6	Basic moderation included.
Set up backups – 2-3 hours.	Included.
24,192,000 Seconds.	60-120 Seconds.

# Cost Efficiencies

- ⦿ No upfront capital expense
- ⦿ Just-in-time capacity management
- ⦿  $10^2$  Cheaper (than what we do now)...



# NEBULA – Designed for the Future

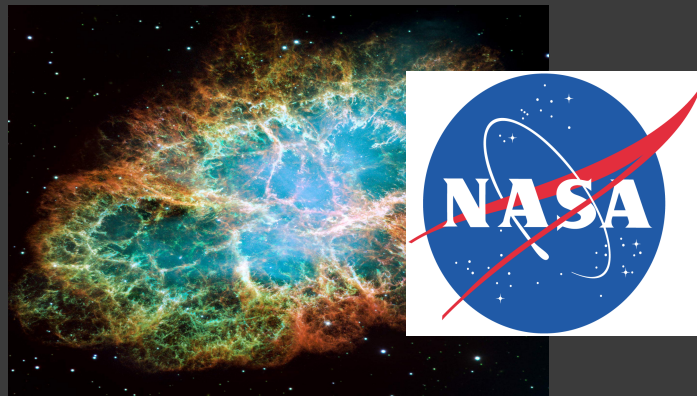
- ◎ “Open standards are critical to the growth of cloud computing and open source software has provided the foundation for many cloud computing implementations.” – Barack Obama
- ◎ NEBULA is 100% Open Source

# NEBULA – the way to Data.gov

- ⦿ Best practices in moderation, open collaboration
- ⦿ Open and Public APIs, everywhere
- ⦿ Feeds (RSS, Atom) power mash-ups
- ⦿ Open-source platform, apps, and data
- ⦿ Fully transparent (we dogfood)

# NASA NEBULA Timeline

- ⦿ Under development since May 2008 (under codename NASA.NET)
- ⦿ Now accepting a limited number of additional partners for Summer, '09
- ⦿ Full Launch, 2010



**WHY  
SUN**  
for cloud  
computing?

- **Innovation  
and  
Choice**

## Sun's Vision



THE NETWORK  
is the  
Computer

Focusing on Network Transparent Computing Since 1984

# Sun's Cloud Computing Strategy

OpenSPARC™



Q-layer



xVM



PostgreSQL



# How Sun Supports Different Approaches

- 
- Public clouds
  - Private clouds (owned and operated)
    - > Customers to run compatible services
    - > Patterns, tools, services, software components, data center design expertise, optimized systems and storage products
  - Partner or hosted clouds
    - > Offer software for service providers and ISVs to reach additional customers and offer compatible variations
  - Common elements
    - > Partnerships, communities and open approach



# Next Steps

# Getting Started

## Developers

- Application inventory
  - Reliability at application or system level?
  - Languages, VMs, tools
- > Developer inventory
  - Internal, communities, partners
- > API inventory
  - Where can new abstractions be inserted?
  - What are deployment mechanisms?

## Data


- Re-factor storage requirements
  - Move storage into the network
  - Structured, semi-structured, relational
- > Analytics foundation
  - All (potential) sources of data
- > Regulatory, security, privacy assessments
  - Access control, audit, assured destruction



# How **SUN** can help

[www.sun.com/cloud](http://www.sun.com/cloud)

- **Cloud Assessment Services**
  - Architectural expertise
  - Alignment of IT infrastructure goals to business driver
  - Industry best practices and proven methodologies to create value-added solutions
- **Cloud Workshops**
  - Gather business requirements
  - Create a high level architecture
- **Specifics**
  - Sun startup essentials
  - Channel partner programs
  - ISV programs
  - Sun learning services
- **Feedback to Sun**
  - Plans in development around a public cloud
  - Open Solaris on EC2



# The Network is ~~the~~ **YOUR** Computer

<http://www.sun.com/cloud>



# Thank You...

**Bill Vass**  
**President/COO**  
**Sun Microsystems Federal, Inc.**

**“Recovering CIO”**  
**[bill.vass@sun.com](mailto:bill.vass@sun.com)**  
**[blogs.sun.com/BVass](http://blogs.sun.com/BVass)**



# Sun Examples

# Comprehensive **OPEN** Portfolio Delivering Customer Choice

Developer Environment	S E R V I C E S		
Database/Storage Platform			
Application Infrastructure			
Virtualization			
Operating System			
Systems Servers Storage Networking			
Microprocessor			

# Sun Virtualization Solutions

- **DESKTOP TO DATACENTER**

innovation – highly integrated to increase efficiencies, security and performance

- **CHOICE**

in virtualization technologies + management of your heterogeneous environments

- Virtualization  
**BUILT-IN**  
and freely available

- **PROVEN EXPERTISE**

to design, implement, optimize and manage enterprise quality dynamic datacenters

# Sun Open Storage

Bringing Simplicity to Cloud Storage

## Open Storage and Archive

- Lower costs with general purpose systems and Open Software
- Seamless integration with existing environments
- Open architectures will free users from vendor lock-in
- Simple data administration
- Storage analytics provide sophisticated, real-time visualization



Dynamic scale with flexible building blocks

Superior performance

Manage more for less

# NY Times: TimesMachine



- Massive data archives
  - > Every newspaper from 1851 to 1922
  - > <http://timesmachine.nytimes.com>
- Utilizes Amazon Web Services (public cloud) and Hadoop (OpenSolaris)
- 405,000 very large TIFF images, 3.3 million articles in SGML and 405,000 xml files -> converted to a more web-friendly 810,000 PNG images and 405,000 JavaScript files
- Created in 36 hours

# Public Cloud Built on Open Source Innovation



- Amazon EC2 expanding customer choice with access to innovative opensource software
  - > Open Solaris – ZFS and Dynamic Tracing (DTrace)
  - > Sun/MySQL – Support for Linux
  - > Java
  - > Hadoop on Open Solaris
- Enhanced options – quick deployment

# Using Sun Technology for Highly Efficient Scaling

## JOYENT

provides Cloud Computing infrastructure and services to Web 2.0 developers and Fortune 500 companies

*Online Infrastructure provider offers hosted storage business with low-cost combined storage and server technology from Sun*

- **Accelerators**

- Virtualized servers deployed within an ecosystem of the highest grade networking and routing fabric available
- Highly scalable on-demand infrastructure for running web sites, including rich web applications written in Ruby on Rails, PHP, Python and Java

- **Partnering with Sun**

- Built using OpenSolaris' Containers Virtualization Technology
- Joyent runs on Sun's CoolThreads and AMD Opteron-based servers running OpenSolaris
- Sun technology allows Joyent to give users and developers access to a great platform at a low price, and offers Web 2.0 developers a powerful environment where they can scale their applications



# TACC: World's Top Supercomputer



- **Sun Fire X4500**
- 72 Systems
- 1.7 petabytes
- 64.8 GB/sec total bandwidth



- **Sun Fire X4600**
- 25 systems
- 800 cores



- **SunBlade 6048**
- 3,936 blades
- 15,744K CPUs
- 62,976 cores
- 125 TB/RAM



- **Sun Data Center Switch 3456**
- Dual redundant
- 110 TB/sec bisectional bandwidth

- The world's largest largest computing system in the world for open science research
- Sun Constellation Linux Cluster and Sun StorageTek Mass Storage Facility
- 579.4 Tflops peak performance